Mellanox MLNX-OS® Release Notes for Ethernet

Software Ver. 3.4.3002

www.mellanox.com

NOTE:
THIS HARDWARE, SOFTWARE OR TEST SUITE PRODUCT ("PRODUCT(S)") AND ITS RELATED DOCUMENTATION ARE PROVIDED BY MELLANOX TECHNOLOGIES "AS-IS" WITH ALL FAULTS OF ANY KIND AND SOLELY FOR THE PURPOSE OF AIDING THE CUSTOMER IN TESTING APPLICATIONS THAT USE THE PRODUCTS IN DESIGNATED SOLUTIONS. THE CUSTOMER'S MANUFACTURING TEST ENVIRONMENT HAS NOT MET THE STANDARDS SET BY MELLANOX TECHNOLOGIES TO FULLY QUALIFY THE PRODUCT(S) AND/OR THE SYSTEM USING IT. THEREFORE, MELLANOX TECHNOLOGIES CANNOT AND DOES NOT GUARANTEE OR WARRANT THAT THE PRODUCTS WILL OPERATE WITH THE HIGHEST QUALITY. ANY EXPRESS OR IMPLIED WARRANTIES, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT ARE DISCLAIMED. IN NO EVENT SHALL MELLANOX BE LIABLE TO CUSTOMER OR ANY THIRD PARTIES FOR ANY DIRECT, INDIRECT, SPECIAL, EXEMPLARY, OR CONSEQUENTIAL DAMAGES OF ANY KIND (INCLUDING, BUT NOT LIMITED TO, PAYMENT FOR PROCUREMENT OF SUBSTITUTE GOODS OR SERVICES; LOSS OF USE, DATA, OR PROFITS; OR BUSINESS INTERRUPTION) HOWEVER CAUSED AND ON ANY THEORY OF LIABILITY, WHETHER IN CONTRACT, STRICT LIABILITY, OR TORT (INCLUDING NEGLIGENCE OR OTHERWISE) ARISING IN ANY WAY FROM THE USE OF THE PRODUCT(S) AND RELATED DOCUMENTATION EVEN IF ADVISED OF THE POSSIBILITY OF SUCH DAMAGE.

Mellanox Technologies
350 Oakmead Parkway Suite 100
Sunnyvale, CA 94085
U.S.A.
www.mellanox.com
Tel: (408) 970-3400
Fax: (408) 970-3403

© Copyright 2015. Mellanox Technologies. All Rights Reserved.

Mellanox®, Mellanox logo, BridgeX®, CloudX logo, Connect-IB®, ConnectX®, CoolBox®, CORE-Direct®, GPUDirect®, InfiniHost®, InfiniScale®, Kotura®, Kotura logo, Mellanox Federal Systems®, Mellanox Open Ethernet®, Mellanox ScalableHPC®, Mellanox Connect Accelerate Outperform logo, Mellanox Virtual Modular Switch®, MetroDX®, MetroX®, MLNX-OS®, Open Ethernet logo, PhyX®, SwitchX®, TestX®, The Generation of Open Ethernet logo, UFM®, Virtual Protocol Interconnect®, Voltaire® and Voltaire logo are registered trademarks of Mellanox Technologies, Ltd.

Accelio™, CyPU™, FPGADirect™, HPC-X™, InfiniBridge™, LinkX™, Mellanox Care™, Mellanox CloudX™, Mellanox Multi-Host™, Mellanox NEO™, Mellanox PeerDirect™, Mellanox Socket Direct™, Mellanox Spectrum™, NVMeDirect™, StPU™, Spectrum logo, Switch-IB™, Unbreakable-Link™ are trademarks of Mellanox Technologies, Ltd.

All other trademarks are property of their respective owners.

Table of Contents

Chapter 1 Introduction 4
Chapter 2 Supported Platforms, Firmware, Cables and Licenses 4
2.1 Supported Switch Systems 4
2.2 Supported CPU Architecture 4
2.3 Supported Firmware 4
2.4 Supported Software Licenses 5
2.5 Upgrade From Previous Releases 5
2.6 Supported Cables 6
Chapter 3 Changes and New Features 7
Chapter 4 Known Issues 12
4.1 General Known Issues 12
4.2 Ethernet Known Issues 15
4.3 IP Routing Known Issues 21
Chapter 5 Bug Fixes 27
5.1 General Bug Fixes 27
5.2 Security Bug Fixes 28
Chapter 6 Submitting a Service Request 31

1 Introduction

This document is the Mellanox MLNX-OS® Release Notes for Ethernet.

MLNX-OS is a comprehensive management software solution that provides optimal performance for cluster computing, enterprise data centers, and cloud computing over Mellanox SwitchX® family. The fabric management capabilities ensure the highest fabric performance while the chassis management ensures the longest switch up time.

The MLNX-OS documentation package includes the following documents:

- User Manual – provides general information about the scope, organization and command line interface of MLNX-OS as well as basic configuration examples
- Release Notes – provides information on the supported platforms, changes and new features, and reports on software known issues as well as bug fixes

2 Supported Platforms, Firmware, Cables and Licenses

2.1 Supported Switch Systems

Table 1 - Supported Switch Systems

Model Number	Description
MSX1036	36-Port QSFP 40GbE 1U Ethernet switch system
MSX1024	12-Port QSFP 40GbE 48-Port SFP+ 10GbE 1U Ethernet switch system
MSX1024(52)	4-Port QSFP 40GbE 48-Port SFP+ 10GbE 1U Ethernet switch system
MSX1016	64-Port SFP+ 10GbE 1U Ethernet switch system
MSX1012	12-Port QSFP (40GbE) 1U Ethernet switch system
MSX1400	48-Port SFP+ 10GbE, 12-Port QSFP 40GbE 1U x86 Ethernet switch system
MSX1410	48-Port SFP+ 10GbE, 12-Port QSFP 40GbE 1U dual core x86 Ethernet switch system (engineering sample)
MSX1710	36-Port QSFP 40GbE 1U Ethernet dual core x86 switch system
MSX1710_O	36-Port QSFP 40GbE 1U Ethernet dual core x86 switch system with ONIE

2.2 Supported CPU Architecture

- PPC 460
- x86

2.3 Supported Firmware

- SwitchX® firmware version 9.3.5080
- SwitchX®-2 firmware version 9.3.5080
- ConnectX®-2 firmware version 2.9.1000 and higher
- ConnectX®-3 firmware version with SwitchX® based systems 2.33.5000 and higher

2.4 Supported Software Licenses

For the software licenses supported with MLNX-OS® software please refer to the “Licenses” section of the “Getting Started” chapter of the *Mellanox MLNX-OS User Manual*.

2.5 Upgrade From Previous Releases

Older versions of MLNX-OS may require upgrading to one or more intermediate versions prior to upgrading to the latest. Missing an intermediate step may lead to errors. Please refer to Table 2 and Table 3 to identify the correct upgrade order.

Table 2 - Supported Software Upgrades for PPC Systems

Target Version	Verified Versions From Which to Upgrade
3.4.3002	3.4.2306; 3.4.2008; 3.4.1120
3.4.2306	3.4.2008; 3.4.1120
3.4.2008	3.4.1954; 3.4.1950; 3.4.1804 ; 3.4.1120; 3.4.0012
3.4.1120	3.4.1110; 3.4.1100; 3.4.0012
3.4.1110	3.4.1100; 3.4.0012
3.4.0012	3.3.5006; 3.3.4402
3.3.5200	3.3.5006; 3.3.4402
3.3.5006	3.3.4402; 3.3.4304
3.3.4402	3.3.4304; 3.3.4100
3.3.4304	3.3.4100; 3.3.3500
3.3.4100	3.3.3500; 3.3.3000

Table 3 - Supported Software Upgrades for x86 Systems

Target Version	Verified Versions From Which to Upgrade
3.4.3002	3.4.2008; 3.4.1120
3.4.2306	3.4.2306; 3.4.2008; 3.4.1120
3.4.2008	3.4.1120; 3.4.0012
3.4.1120	3.4.1110; 3.4.1100; 3.4.0012; 3.3.5006
3.4.1110	3.4.1100; 3.4.0012; 3.3.5006
3.4.0012	3.3.5006
3.3.5200	3.3.5006; 3.3.4402
3.3.5006	3.3.4406; 3.3.4304

For upgrade instructions refer to the section “Upgrading MLNX-OS Software” in *Mellanox MLNX-OS User Manual*.

Downgrading from MLNX-OS version 3.4.0012 or later on x86 switch systems is not supported.

2.6 Supported Cables

For a list of the Mellanox supported cables please visit the LinkX™ Cables and Transceivers page of the Mellanox Website at http://www.mellanox.com/page/cables?mtag=cable_overview.

When using Mellanox AOC cables longer than 50m use one VL to achieve full wire speed.

3 Changes and New Features

Table 4 - Ethernet Changes and New Features

Category	Description
Release 3.4.3002	
Security	Added support for switch secure mode
User Accounts	Improved logic of AAA authorization map order See the command “aaa authorization map order” in the User Manual
System Management	Added support for auto-BIOS upgrade for SX1400 and SX1710 switch systems
Chassis Management	Added new System Health Monitor notification alert See the section “System Health Monitor Alerts Scenarios” in the User Manual
CLI	Improved module status display See command “show module” in the User Manual
Software Management	Added support for MLNX-OS over ONIE See section “Image Maintenance via Mellanox ONIE” in the User Manual
Software Management	Removed requirement for 56GbE license See section “Licenses” in the User Manual
XML API	Improved XML interface Refer to MLNX-OS® XML API Reference Guide for more information
Ethernet Switching	Added support for QinQ See section “QinQ” in the User Manual
802.1x Protocol	Added support for single-host 802.1x authentication protocol See section “802.1x Protocol” in the User Manual
PIM	Added support for next-hop load balancing See section “Next Hop Load-Sharing” in the User Manual
Release 3.4.2008	
CLI	Modified user interface for the command “show asic version”
CLI	Improved module hierarchy in the output of the commands “show power” and “show temperature”
CLI	Removed CPU component from the output of the command “show inventory”
SNMP	Applied new index scheme for SNMP EntityTable
Ethernet Switching	Added support for multiport 802.1x authentication protocol See section “802.1x Protocol” in the User Manual
IP Routing	Added support for VRF with static IPv4 and ECMP See section “Virtual Routing and Forwarding” in the User Manual
PIM	Added GA support for PIM Load-Sharing See section “PIM Load-Sharing” in the User Manual

Table 4 - Ethernet Changes and New Features

Category	Description
IPv6	Added support for IPv6 over loopback port
MLAG	Support for MLAG fast recovery
Release 3.4.1120	
General	Removed “sx_” prefix from version numbers in the code
Release 3.4.1110	
General	Bug fixes
Release 3.4.1100	
LACP	Added support for specific port in the command “show lacp neighbor”
Ethernet Switching	Added GA support for RPVST See section “RPVST” in the User Manual
Ethernet Switching	Added support for LACP individual mode See the command “suspend-individual” in the User Manual
Ethernet Switching	OpenFlow bitwise IP match
Ethernet Switching	Added support for BPDU guard See the command “spanning-tree port type” in the User Manual
IP Routing	Added GA support for IPv6 See section “IPv6 Routing” in the User Manual
IP Routing	Added support for VRRP counters
IP Routing	Added prefix filter in the command “show ip route <ip-address>”
MAGP	Added “MAGP disabled” indication in the command “show magp” when applicable
SNMP	Added support for SNMP SET upgrade See section “Upgrading MLNX-OS Software with SNMP” in the User Manual
WebUI	Added popup Welcome screen when connecting via WebUI See section “Starting the Web User Interface” in the User Manual
Security	Added default passwords to the XML default users See section “User Accounts” in the User Manual
Release 3.4.0012	
Security	Changed the HTTPS default ciphers to TLS.
IP Routing	GA support for PIM Bi-Dir. See section 6.3 “Multicast (IGMP and PIM)” in the User Manual.
IP Routing	Added support for ICMP redirect. See command “ip icmp redirect” in the User Manual.
Ethernet Switching	Added support for disabling MAC-address learning.

Table 4 - Ethernet Changes and New Features

Category	Description
General	Added support for Mellanox OFED 2.3 integration.
Interconnect	Added support for LR4 modules on Ethernet switch systems.
IP Routing	GA support for BGP. See section 6.3 “BGP” in the User Manual.
MLAG	Added support for MLAG LACP.
Security	Added support for security strict mode.
SNMP	Added support for setting hostname through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
SNMP	Added support for power cycling through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
SNMP	Added support for changing configuration through SNMP. See section 4.17.1.7 “SNMP SET Operations” in the User Manual.
Systems	GA support for SX1710 switch systems.
Virtual Machine	GA support for Virtual Machine for x86-based switch systems. See section 4.19 “Virtual Machine” in the User Manual.
WebUI	Added support for Internet Explorer 11 web browser.
Release 3.3.5200	
IP Routing	GA support for Router Port. See Chapter 6 “IP Routing” in the User Manual.
IP Routing	Proof of concept support for BGP. See section 6.3 “BGP” in the User Manual.
IP Routing	GA support for route maps. See section 6.4 “Route Map” in the User Manual.
IP Routing	GA support for PIM Bi-Dir. See section 6.6 “Multicast (IGMP and PIM)” in the User Manual.
IGMP Querier	GA support for IGMP Querier. See section 5.8.3 “IGMP Snooping Querier” in the User Manual.
SNMP	Added SNMP traps for VLAN Interface, MLAG port channel and Port channel link up/down events.
SNMP	Added MSTP MIBs/traps support.
Release 3.3.5006	
Ethernet Switching	GA support for MLAG.
ACLs	Added MAC-based VLAN ACL.
IP Routing	GA support for MAGP.

Table 4 - Ethernet Changes and New Features

Category	Description
IP Routing	GA support for VRRP.
IP Routing	GA support for IGMP.
IP Routing	Alpha support for PIM.
OSPF	Loopback interface on OSPF.
OSPF	Added new OSPF MIB/Traps.
Release 3.3.4664	
IP Routing	Alpha support for MAGP.
IP Routing	Alpha support for VRRP.
SNMP	Added cable info entries to entPhysicalTable.
SNMP	Added support for SNMP to trigger SNMP test trap via SNMP set command. See section “MLNX-EFM MIB”.
SNMP	Added system identifier (MAC address) to test trap.
Release 3.3.4402	
General	Added new certificate hashing algorithm (sha256). See section “Cryptographic (X.509, IPSec)” in the User Manual.
General	The command “show configuration full” is no longer supported.
CLI	Added support for command output filtering. See section “Command Output Filtering” in the User Manual.
DCBX	Added support for LLDP-DCBX CEE TLVs.
Release 3.3.4304	
General	Added End-User License Agreement. See section “Getting Started” in the User Manual.
General	Improved configuration file format.
Power Management	Added support for link width reduction.
Systems	New CPU temperature for MSX10xx systems.
Systems	Beta support for x86 CPU family-based systems (MSX1400 and MSX1700).
Puppet Agent	GA support for Puppet. See section “Puppet Agent” in the User Manual.
Ethernet Switching	GA support for debug per protocol.
OpenFlow	GA support for OpenFlow. See section “OpenFlow” in the User Manual.
DCBX	GA support for DCBX. See section “DCBX” in the User Manual.

Table 4 - Ethernet Changes and New Features

Category	Description
MSTP	GA support for MSTP protocol. See section “MSTP” in the User Manual.
DHCP Relay	GA support for DHCP relay. See section “DHCP Relay” in the User Manual.
IGMP Snooping	GA support for IGMP Querier relay. See section “IGMP Snooping Querier” in the User Manual.
Release 3.3.4100	
General	Improved debug file upload mechanism. Refer to “file debug-dump” command in the CLI reference guide.
General	Added support for displaying system hardware revision. Refer to “show inventory” command in the CLI reference guide.
SNMP	Added a MELLANOX-ENTITY-MIB as an extension to the standard ENTITY-MIB to represent system GUID. Refer to SNMP chapter in the User Manual. The new MIB can be found in Mellanox support website.
Logging	Added support for event notification to monitor. Refer to “Event Notification” section in the User Manual.
Chassis Management	Improved temperature control algorithm.
User Interfaces	Improved login timeout mechanism.
OSPF	NSSA area type support. Refer to “area nssa” command in the OSPF chapter of the CLI guide.
Release 3.3.4000	
Systems	SX1012 supported at GA level.
General	Bug fixes.
Release 3.3.3500	
IP Routing	OSPF support (GA level).
Management Interfaces	In-band management.
Ethernet Switching	Port mirroring.
Ethernet Switching	sFlow.
IP Routing	OSPF support (beta level).

4 Known Issues

The following sections describe MLNX-OS® known issues in this software release and possible workarounds.

For hardware issues, please refer to the switch support product page.

4.1 General Known Issues

Table 5 - General Known Issues (Sheet 1 of 4)

Index	Category	Description	Workaround
1.	Management Interfaces	The command `reset factory keep-basic` removes management IP configuration.	N/A
2.	Management Interfaces	The CLI command `ip default-gateway <interface>` sets the gateway address to 0.0.0.0 and prevents the user from adding other gateways.	Delete the entry by using the command `no ip default-gateway`.
3.	Management Interfaces	The “no” negation prefix is not supported prefix in the following CLI commands: • interface ethernet <s/p> ipv4 port access-group nodhcprelay • interface port-channel <num> switchport mode	Use the “no” keyword as an infix (e.g. `interface port-channel <num> no switchport`).
4.	Management Interfaces	Switch systems may have an expired HTTPS certification.	Generate a new certificate by changing the hostname.
5.	Management Interfaces	Consecutive hostname modification is not supported.	Wait 25 seconds before reattempting to modify the hostname.
6.	Puppet Agent	Only the resource ID “default” is supported for Puppet router.	N/A
7.	Puppet Agent	Duplex parameter is not supported for Ethernet or LAG interfaces.	N/A
8.	In-Band Management	In-band management mode does not support IPv6.	N/A
9.	In-Band Management	In-band management ARPs are not taken into account in the command `show ip arp count`.	N/A
10.	In-Band Management	In-band management is not operational without an L3 license.	Contact Mellanox to get L3 license for free.
11.	NTP	The command `show ntp` always lists the last configured NTP server even if it has been deleted. This output can be safely ignored.	N/A

Table 5 - General Known Issues (Sheet 2 of 4)

Index	Category	Description	Workaround
12.	Software Management	Only one image is allowed to be copied into the system (using the `image fetch` command). The user must remove old image files prior to fetching a new one.	N/A
13.	User Accounts	If AAA authorization order policy is configured to remote-only, then when upgrading to 3.4.3002 or later from an older MLNX-OS version, this policy is changed to remote-first.	N/A
14.	Configuration Management	Loading the wrong system configuration file may hang the system. For example: an attempt to use an SX1016 configuration file on an SX6036/SX6536/SX1035 system may cause the system to hang.	N/A
15.	Configuration Management	After loading a new configuration file, please reboot the system. Otherwise, configuration may not be properly applied and errors may appear in the log.	N/A
16.	Configuration Management	When using a large set of configuration files, configuration apply can take more time than usual due to parallel activity of statistics data collecting.	N/A
17.	Configuration Management	Applying a configuration file of one system profile to another is not supported.	N/A
18.	Configuration Management	If the last box leaving an HA (MLAG/) cluster is not the box on which the cluster was created, bogus configuration will result which does not allow recreating the same cluster name.	N/A
19.	Configuration Management	Sending packets to a non-default port in TFTP transport layer is not supported.	N/A
20.	Logging	“DROPPED MSG” errors may appear during reload (shutdown phase). These errors can be safely ignored.	N/A
21.	Logging	The warning “pgm_set_timeout” may appear in the log. This warning can be safely ignored.	N/A
22.	Logging	During system de-init, the error “[mdreq.ERR]: init(), mdr_main.c:634, build 1: Error code 14014” may appear in the log. This error can be safely ignored.	N/A

Table 5 - General Known Issues (Sheet 3 of 4)

Index	Category	Description	Workaround
23.	Logging	The warning "[mgmtd.WARNING]: Upgrade could not find node to delete: /iss/config/stp/switch/ethernet-default/spanning-tree/mode" may appear in the log. This warning can be safely ignored.	N/A
24.	Logging	When using a regular expression containing | (OR) with the command `show log [not] matching <reg-exp>]`, the expression should be surrounded by quotes ("<expression>"), otherwise it is parsed as filter (PIPE) command.	N/A
25.	Logging	Port up/down events on a port quickly toggling states may be displayed in wrong order in the monitoring terminal.	For actual port stats, use the command `show interface`.
26.	Logging	The following error may appear in the log: "sn_mdc_msg_handler(), ../../tms/sn_mgmt.c:910, build 1: Error code 14002 (assertion failed) returned". This error may be safely ignored.	N/A
27.	User Management	Some RADIUS and TACACS+ configurations keep the user locked out of the machine due to timeout limitation.	Press the reset button for 15 seconds, and then log in using your local authentication. Additionally, fix the configuration to avoid any future timeout issues.
28.	WebUI	Reversing the time clock can result in WebUI graphs' corrupted data.	Clear the graphs data after setting the clock.
29.	WebUI	Enabling/disabling HTTPS while connected via HTTP to the WebUI may result in temporary loss of connection to the webpage.	Refresh the page or navigate back using the browser's back button.
30.	WebUI	The following commands cannot be executed from "Execute CLI commands" section on WebUI. • interface ethernet <interface> module-type • dcb priority-flow-control enable	To run those commands use the CLI.
31.	WebUI	If the configured ciphers in versions prior to 3.4.0012 were SSL and TLS ciphers, upgrading to this version will override that. The new default is to allow TLS ciphers only. To enable SSL, please run the command `web https ssl ciphers all`.	N/A
32.	WebUI	When SSH strict mode is activated with TLS 1.2, Firefox does not work properly.	N/A

Table 5 - General Known Issues (Sheet 4 of 4)

Index	Category	Description	Workaround
33.	WebUI	When upgrading to version 3.4.3002, statistics files are reset. As a result, WebUI statistic graphs are reset as well.	N/A
34.	CLI	MLNX-OS support up to 50 CLI session open in parallel.	N/A
35.	CLI	Command output filtering does not support the following commands: • show log • show puppet-agent log • show configuration text files <file>	N/A
36.	SNMP	The error “Cannot find module (MELLANOX-MIB)” may appear in the log when performing rollback to a MLNX-OS version older than 3.3.3000. This error can be safely ignored.	N/A
37.	SNMP	Upon system shutdown, the following error may appear: “[mibd.ERR]: mdc_foreach_binding_ prequeried_parsed(), mdc_main.c”. This error can be safely ignored.	N/A
38.	SNMP	The ifNumbers MIB (OID: 1.3.6.1.2.1.2.1.0) on x86 switch systems displays 42 interfaces while the ifTable displays 40 due to VM management interfaces that are not shown in the ifTable.	N/A
39.	Chassis Management	Duplicate notification is sent upon high temperature alerts, both specific module alerts and general “too high” alerts.	N/A
40.	Chassis Management	When plugging in the power supply without connecting it to the power source, the system status and power supply LEDs become RED.	Plug in the power source.
41.	Virtual Machine	For volume fetch, using a USB drive formatted with VFAT causes errors in the log and may require additional reboot for the USB to be registered for virtual machine volume usage.	Use EXT3 USB format.

4.2 Ethernet Known Issues

Table 6 - Ethernet Known Issues (Sheet 1 of 7)

Index	Category	Description	Possible Workaround
1.	Ethernet Interfaces	The system allows tolerance of 4 bytes on top of the set MTU.	N/A
2.	Ethernet Interfaces	In case of faulty cable, rise time of other interfaces in the system may be delayed.	N/A

Table 6 - Ethernet Known Issues (Sheet 2 of 7)

Index	Category	Description	Possible Workaround
3.	Ethernet Interfaces	40GbE passive copper cable of 5m length might experience link rise issues when connected to some 3rd party 40GbE interfaces.	Replace copper cable with fiber cable
4.	Ethernet Interfaces	Default MTU is modified as 1500 instead of 1522.	N/A
5.	Ethernet Interfaces	There are no port counters for packet sizes 1518-1522.	N/A
6.	Ethernet Interfaces	Traffic loss may occur when enabling flow control on a packet with an MTU larger then 8K on SX1012.	N/A
7.	Ethernet Interfaces	LAG cost is not calculated correctly.	N/A
8.	Ethernet Interfaces	56Gb/s speed is not supported on SwitchX® (A1) ASIC.	N/A
9.	Ethernet Interfaces	If using eth-single-swid system profile, static ARP cannot be configured on the mgmt0 interface without an L3 license.	N/A
10.	User Interfaces	SX1012 hybrid cable information may not appear on WebUI/CLI.	N/A
11.	LAG	Convergence time of port removal from a LAG exceeds the specification by up to 30 seconds.	Shut down the port prior of removing it from the LAG.
12.	LAG	Traffic running over LAG may not be evenly distributed when testing small number of SMACs.	N/A
13.	LAG	When converting static LAG to dynamic LAG (LACP), some of the ports remain in DOWN state.	Disable and enable the LAG.
14.	LAG	LAG Hash function lacks configuration of a randomization seed parameter. Traffic passed over LAG in hierarchy topology via multiple SwitchX based systems is not distributed evenly.	Change the distribution function between layers performing LAG using the command `port-channel load-balance`. The workaround is applicable in case there are multiple mixed flows.
15.	LAG	When removing all ports of a port-channel and then re-adding them to the port-channel in a different type, port-channel may not converge.	Disable and enable the port-channel.

Table 6 - Ethernet Known Issues (Sheet 3 of 7)

Index	Category	Description	Possible Workaround
16.	LAG	The following LACP port-states are not available in MLNX-OS®: • I – Individual H – hot-standby (LACP only) • s – Suspended R – module-removed	N/A
17.	LAG	Configuring multiple LAGs as a range causes the switch to hang for up to 10 minutes.	N/A
18.	MLAG	Each MLAG VIP group must be configured with a different unicast IP address. If not, MLAG behavior is not anticipated.	N/A
19.	MLAG	MLAG is not supported in WebUI.	N/A
20.	MLAG	MLAG port-channels (Mpo) appear as regular port-channels (Po) in WebUI.	N/A
21.	MLAG	MLAG configuration including (mlag-port-channel configuration) must be identical in both switches.	N/A
22.	MLAG	MLAG VIP is limited to 2 switches. If more than 2 switches connected to the same MLAG VIP, MLAG behavior is not anticipated.	N/A
23.	MLAG	When the MPO is connected to the host, spanning tree should be disabled on the MLAG port-channel before it is enabled (port type set to edge and BPDU filter enabled). When connected to a switch, spanning tree should be disabled on the switch (`no spanning-tree`).	N/A
24.	MLAG	Deleting an IPL port-channel interface while MLAG is enabled (using the command `no interface port-channel <id>`) might cause errors. These errors may be safely ignored.	Shutdown the port before deleting it.
25.	MLAG	The IPL VLAN interface must be used only for MLAG protocol and must not be used by any other interfaces (e.g. port-channel, Ethernet).	N/A
26.	MLAG	An MLAG-port-channel must be administratively disabled (`interface mpo <id> shutdown`) before deleting it.	N/A
27.	MLAG	IGMP snooping is automatically enabled on MLAGs and may not be disabled.	N/A
28.	MLAG	Keepalive interval should be set to 3 seconds when working with over 1k VLANs on PPC.	N/A

Table 6 - Ethernet Known Issues (Sheet 4 of 7)

Index	Category	Description	Possible Workaround
29.	MLAG	Configuring a LAG to become an IPL when another LAG is already configured as IPL leads to MLAG going down and an error being printed in the log.	Have only one LAG set as IPL. Remove any previous configuration if a new one is to be applied.
30.	MLAG	Changing STP mode affects traffic and may cause MLAG to stop working for an interval of 21 seconds until the new STP topology is built.	User should disable MLAG prior to changing the STP mode.
31.	MLAG	Adding and removing VLAN ranges on trunk/ hybrid ports is correlated to the number of the trunk/hybrid ports in the system.	Restrict the range of VLANs added/removed to/from an MLAG port to groups of up to a 500 VLANs.
32.	VLAN	It might take approximately 20 seconds to create/delete 1000 VLANs, or to change the switchport mode of an interface.	N/A
33.	VLAN	Creating/deleting VLAN ranges on trunk/ hybrid ports is correlated to the number of the trunk/hybrid ports in the system.	Restrict the range of VLANs created/deleted to groups of up to a 70 VLANs.
34.	VLAN	When working with more than 2000 VLANs and trying to change a port's VLAN membership in modes hybrid/trunk using the command `interface 1/x switchport trunk allowed-vlan`, the operation may timeout on account of it being too heavy. Thus, the port does not become properly configured as a member of those VLANs. Assuming port 1 is in trunk mode and there are 3000 VLANs created, a good practice would be to achieve membership of port 1 in VLANs 100-3000, use the command `interface ethernet 1/x switchport trunk allowed-vlan remove 1-100` instead of `interface ethernet 1/x switchport trunk allowed-vlan 100-3000`.	To fix the problem, run the commands `interface ethernet 1/x switchport trunk allowed-vlan all` and `interface ethernet 1/x switchport trunk allowed-vlan remove 1-100` in succession.
35.	Spanning Tree	Non-standard behavior may occur on third party switch STP functionality when the port moves from Root to Alternate state. The standard Proposal-Agreement sent by MLNX-OS is unanswered which may result in traffic loss of up to 30 seconds.	N/A
36.	Spanning Tree	STP convergence time is slower than expected.	N/A
37.	Spanning Tree	BPDUs are sent from both the superior and the inferior switches.	N/A

Table 6 - Ethernet Known Issues (Sheet 5 of 7)

Index	Category	Description	Possible Workaround
38.	MSTP	The command `no vlan` on a VLAN interface associated with an MSTP instance is not supported.	Dissociate the VLAN from the MSTP instance before running the command.
39.	MSTP	Under heavy MSTP configuration, the system may exhibit slowness after adding a new MSTP instance. In the ensuing minute, if a show command is issued an error message may appear. That message can be safely ignored.	N/A
40.	IGMP Snooping	Setting the `ip igmp snooping unregistered multicast` option to `forward-to-mrouter-ports` option reduces the amount of supported VLANs to 250.	N/A
41.	IGMP Snooping	There is no possibility to have more than one mrouter port on a single VLAN if IGMP snooping is disabled for that VLAN.	After enabling IGMP snooping on that VLAN, the user needs to re-add the mrouter-ports on that VLAN.
42.	IGMP Snooping	When no IP interfaces exist on the switch, the default IGMP querier address is 10.0.0.1.	Create an interface VLAN with the required IP address and reconfigure the IGMP querier on the required VLAN.
43.	IGMP Snooping	IGMP snooping is limited to a single MC IP per MC MAC.	N/A
44.	LLDP	LLDP notifications are not supported.	N/A
45.	LLDP	The mgmt0, mgmt1 IPv6 addresses and the management address are not advertised on LLDP TLV(s).	N/A
46.	LLDP	When configuring LLDP with PFC on a LAG, after switch reload, the LLDP does not show the PFC configuration anymore.	After switch reload, reconfigure the PFC LLDP on the LAG using the command `interface ethernet <x/y/z> lldp tlv-select dcbx`.
47.	LLDP	LLDP ifIndexes do not match the ones published in the ifTable in IEEE 802.1.	Use the ifDescr and fields.
48.	ACLs	Packets dropped by the switch due to congestion or ACL rules are added to the bad type counter of the matching VLAN interface.	N/A
49.	ACLs	When IP Routing is enabled, the maximum number of ACLs that can be configured is 3	N/A

Table 6 - Ethernet Known Issues (Sheet 6 of 7)

Index	Category	Description	Possible Workaround
50.	Modules	When using 5m splitter cable (P/N MC2609125-005), wrong data is displayed under transceiver information.	N/A
51.	Modules	Using 1GbE copper modules leads to the port status being constantly reported as active regardless from the actual link state.	N/A
52.	Modules	The command `show interface` does not display 1GbE modules correctly.	N/A
53.	SNMP	LLDP MIB OID lldpRemPortDesc is not supported.	N/A
54.	Puppet Agent	Speed parameter is not supported for LAG interfaces.	N/A
55.	Puppet Agent	Description parameter is not supported on L2 interfaces.	N/A
56.	Port Mirroring	When port mirroring congestion is configured to “drop-excessive-frames” (best effort), while using the same analyzer port for more than a single mirroring session, that port could drops packets from all flows not according to their configured priorities.	N/A
57.	sFlow	The discarded packets counter in sFlow samples may not be accurate and may expose a skew of one second.	N/A
58.	sFlow	Discard counters are not reported to the collector.	N/A
59.	sFlow	Errors are seen when defining sFlow session from two WebUI or CLI sessions simultaneously.	N/A
60.	OpenFlow	OpenFlow ARP packet-matching rule does not include IP source and destination addresses.	N/A
61.	OpenFlow	OpenFlow is not supported over LAG, member ports of LAG, or split-port interfaces.	N/A
62.	OpenFlow	OpenFlow can support up to 1000 flows, supporting 12 tuples with wildcards according to the standard.	N/A
63.	OpenFlow	Standard ACL cannot be configured while OpenFlow is in use.	N/A
64.	OpenFlow	OpenFlow counters can support counting of packets.	N/A

Table 6 - Ethernet Known Issues (Sheet 7 of 7)

Index	Category	Description	Possible Workaround
65.	OpenFlow	OpenFlow packet modify action supports VLAN (setting VLAN ID, setting VLAN priority and stripping VLAN header – for QinQ packets).	N/A
66.	OpenFlow	Output can only be set to `flood`, `physical port normal`, `controller` and `drop`.	N/A
67.	OpenFlow	OpenFlow EtherType-matching rule can be set to one of: IPv4, IPv6, ARP/RARP, FCoE and 2 user defined Ethertypes.	N/A
68.	OpenFlow	OpenFlow ICMP packet-matching rule does not include type and code.	N/A
69.	DCBX	Enabling LLDP triggers a faulty notation of PFC oper-state. This status should be ignored.	N/A
70.	DCBX	There is no automatic fallback between IEEE and CEE.	N/A
71.	DCBX	When moving from DCBX CEE to DCBX IEEE, TLVs may be sent inappropriately.	After mode change, the TLVs to be sent must be reset.
72.	DCBX	DCBX auto-select type (IEEE/CEE) is not supported.	DCBX type should be selected manually.
73.	Chassis Management	Problems with the power supply cannot be monitored on SX1016 switch systems using the command `show health-report` or by observing the PS LED on the switch system itself.	N/A
74.	802.1x	Show dot1x commands on PPC based systems provide output in a slow line-by-line fashion.	N/A
75.	802.1x	A PAE may send unneeded EAP-START messages when authentication starts.	N/A

4.3 IP Routing Known Issues

Table 7 - IP Routing Known Issues (Sheet 1 of 6)

Index	Category	Description	Possible Workaround
1.	IP Interfaces	ICMP echo packets are not counted as part of the VLAN interface counters.	ICMP packets are counted in the interface L2 counters.
2.	IP Interfaces	The maximum number of L3 interfaces that can be configured is 64	N/A

Table 7 - IP Routing Known Issues (Sheet 2 of 6)

Index	Category	Description	Possible Workaround
3.	IP Interfaces	A VLAN bound with an ACL action must not be bound to a router port. Else, router port behavior is unpredictable.	N/A
4.	IP Interfaces	Router ports are not supported on SX10xx-xxxR and SX60xx-xxxR systems.	N/A
5.	IP Interfaces	It is not possible to add or remove member interfaces from a LAG, or delete that LAG when it is configured as a router port.	Physical port must explicitly be changed from "no switchport" to "switchport" so that router capabilities are reset from physical port.
6.	IP Interfaces	ICMP redirect is not supported over router port.	N/A
7.	LLDP	A port connected to a router-port receives from it an LLDP packet containing IPv4 and IPv6 management addresses. However, the IPv6 address is not included in the LLDP MIB as a remote management address. Only the IPv4 remote management address is published.	N/A
8.	IP Diagnostic Tools	While using invoked shell commands such as ping, traceroute, tcpdump, for egressing from a specific router interface use `-I/i vlan<vlan id>`. For example, to ping 10.10.10.10 via interface VLAN 30 run the command: `ping 10.10.10.10 -I vlan30`.	N/A
9.	IP Routing	Prior to upgrading to this software release the user must align the number of configured multicast routes to the number of routes defined in section "IPv4 Routing Mode" in the User Manual.	N/A
10.	IP Routing	LAG and L3 ECMP Hash functions lack the configuration of a randomization seed parameter. Traffic passed over LAG or L3 ECMP in hierarchy topology via multiple SwitchX based systems is not distributed evenly.	Change the distribution function between layers performing LAG or L3 ECMP using the command `port-channel load-balance / ip load-sharing`. The workaround is applicable in case there are multiple mixed flows.

Table 7 - IP Routing Known Issues (Sheet 3 of 6)

Index	Category	Description	Possible Workaround
11.	IP Routing	In IP Routing mode, multicast or broadcast packets such as ARP may be counted as "RX discard" in case only one interface is a member of the VLAN. These kind of packets do reach the CPU for processing. The progressing counter can be ignored.	N/A
12.	IP Routing	Errors in the log may appear when using ECMP routes if next-hop belongs to a VLAN interface which lies on port-channel whose links operational state is toggling.	N/A
13.	IP Routing	If there is a static ARP on top of an interface VLAN and ECMP route with its next-hop equal to the static ARP entry, and if the static ARP owner is not responding to ARP requests, then performing `shutdown` and `no shutdown` to the interface VLAN, within a minute, causes errors to appear on the log. These errors may be safely ignored.	N/A
14.	IP Routing	ARP aging timeout minimum value is 240. After upgrade to release 3.3.3500, any value lower than this is clamped to 240.	N/A
15.	IP Routing	The configuration of `ip arp timeout` does not affect entries related to gateway, out-of-band, or management interfaces.	N/A
16.	IP Routing	If the number ARP entries configured on a system is over the permitted limit, the machine becomes overloaded and hangs.	Power cycle the system.
17.	IP Routing	If static route next-hop of a VLAN interface is configured before configuring its subnet, that configured route is not added to the interface and cannot be deleted later.	Make sure to configure the VLAN interface subnet before configuring the next-hop itself.
18.	IP Routing	Global and interface-specific ARP timeout configuration commands do not function.	N/A
19.	IP Routing	The command `ip arp timeout` does not take effect on L3 interfaces which have been created before changing the ARP timeout.	Reload the system after running the command.
20.	OSPF	Changing the IP address on an OSPF Interface VLAN deletes the interface-OSPF attachment.	Reconfigure all OSPF settings after changing the IP.
21.	OSPF	Interface configuration cannot be performed while OSPF is admin-disabled (shutdown).	Enable OSPF admin state using the command `router OSPF no shutdown`.

Table 7 - IP Routing Known Issues (Sheet 4 of 6)

Index	Category	Description	Possible Workaround
22.	OSPF	Removing a static ID using the command `no router-id` requires at list one VLAN interface to be configured.	N/A
23.	OSPF	OSPF supports learning of up to 4K routes.	N/A
24.	OSPF	Default route advertise by ABR to NSSA area not installed by NSSA area router.	Install the default route manually on the NSSA area using the command `ip route`.
25.	OSPF	Changing administrative distance (the command `router ospf distance`) is not supported.	N/A
26.	OSPF	Reverse packet filtering is not supported. When misconfigured, traffic may pass on an unintended subnets.	N/A
27.	OSPF	If the MTU is reconfigured while the interface is up, disruptions may be experienced in L3 traffic.	Shut down the interface before reconfiguring MTU.
28.	OSPF	A loopback interface assigned to area 0 does not get redistributed as direct route after removal from the OSPF.	Keep the loopback assigned to the OSPF area 0
29.	DHCP Relay	DHCP Relay is supported on SwitchX-2 platforms only.	N/A
30.	DHCP Relay	DHCP Relay is not supported on SX10xx-xxxR and SX60xx-xxxR systems.	N/A
31.	VRRP	Using the same Virtual IP (VIP) as the VLAN interface is not supported.	N/A
32.	MAGP	Designated traffic, such as ping, to the MAGP interface is not supported.	N/A
33.	PIM	Connecting two routers via VLAN interface with PIM load-sharing causes loops in the network.	Use a router port instead
34.	PIM	PIM BIDIR routers must have point-to-point connection.	N/A
35.	PIM	The command `show ip pim rp` does not distinguish elected RP from other active RPs.	Use the command `show ip pim rp-hash` to see the elected RP, and `show ip mroute` to see active RPF-interface.

Table 7 - IP Routing Known Issues (Sheet 5 of 6)

Index	Category	Description	Possible Workaround
36.	PIM	Per-interface multicast TTL threshold, and static OIF cannot be configured without enabling PIM in the interface.	Enable PIM-SM on the interface.
37.	PIM	Local setting of candidate-BSR interval affects the BSR hold-time of other winner BSR router.	Use the default BSR interval on the BSR router (60 seconds), or configure all routers with the same interval. (Note that this does not affect per-RP hold-time within BSR).
38.	PIM	At most 300 IGMP packets can be processed per second.	Do not use PIM in environments where there are more than 300 joins/ leaves per second.
39.	PIM	RP candidate can be enabled only on one interface at a time.	N/A
40.	PIM	BSR candidate can be enabled only on one interface at a time.	N/A
41.	PIM	In an L3 network with two or more PIM-BIDIR routers, IGMP snooping should be disabled.	N/A
42.	PIM	Updating BSR hash mask only affects new multicast groups, while groups already active still use the old map. This may lead to loops in the network.	Reboot all routers in the network.
43.	Multicast	Multicast protocol (PIM) must be enabled in order to configure static multicast route.	Enable PIM.
44.	Multicast	Router port does not filter outgoing multicast packets according to configured multicast time-to-live (TTL) value.	N/A
45.	BGP	Route aggregation configuration cannot be modified.	Delete route aggregation and create a new configuration.
46.	BGP	Aggregate overlapping not supported	N/A
47.	BGP	Using the command `clear ip bgp <ip> soft in` mandates enabling route-fresh on the peer.	If route-refresh is not supported on the peer device, use the command `clear ip bgp <ip> in` instead.
48.	BGP	The command `bgp aggregate-address summary-only` does not silence aggregate-route re-advertisements.	N/A

Table 7 - IP Routing Known Issues (Sheet 6 of 6)

Index	Category	Description	Possible Workaround
49.	BGP	AS-path request/set are limited to 128. When more are received, TCP FIN is sent and connection to the peer is lost.	N/A
50.	BGP	When running both BGP and OSPF, changing the configuration of the command `router bgp <AS> <external distance> <internal distance> <local distance>` requires running `clear ip bgp all` afterwards.	N/A
51.	BGP	Four-byte ASN is not supported.	N/A
52.	BGP	The command `no neighbor <ip-address> local-as` is not functional.	Delete the neighbor with `no neighbor {<ip-address>} remote-as` and restore the neighbor without the command `neighbor local-as`.
53.	ACL	When upgrading to this release, it is advised to reduce the number of ACL configured to 3.	N/A
54.	VRF	A router port LAG cannot be mapped to user VRF.	N/A
55.	VRF	Any VRF can be monitored with sFlow, but a collector can be on the management VRF only.	N/A

5 Bug Fixes

5.1 General Bug Fixes

The following table describes MLNX-OS® bug fixes in this software release.

Table 8 - General Bug Fixes

Index	Category	Description
1.	SNMP	SNMP EntityTable does not refresh immediately after an event.
2.	Chassis Management	The command “show module” displays incorrect “Power” status.
3.	Chassis Management	On rare occasions when sending MAD queries which require I^2C access (e.g. cable information) to SX1710 switch systems, the transaction may hang.
4.	Chassis Management	Over-temperature alert on QSFP module is sent on wrong threshold.
5.	Chassis Management	Power supply may appear as not present but powered when extracted from the system.
6.	Chassis Management	The internalSpeedMismatch trap is raised only once the first time it is discovered.
7.	Chassis Management	Syslog messages, and trap messages for internal link speed mismatch do not include the ifName of the link which identify both ends of the connection.
8.	User Accounts	Setting AAA authorization mapping to remote-only does not work. Local credentials are still used.
9.	User Accounts	ASCII based authentication using TACACS+ is not functional.
10.	System Management	Received SysRq signals from serial connection (RS232) to USB adapter can cause switch to reboot.
11.	System Management	On rare occasion, MGMT1 link may go down on SX1710 switch systems.
12.	XML API	XML node for “interface vlan enable/disable” does not work.

5.2 Security Bug Fixes

Table 9 presents the security bug fixes which are added in this MLNX-OS version.

Table 9 - List of Security Bug Fixes

CVE	Description
CVE-2013-7423	The send_dg function in resolv/res_send.c in GNU C Library (aka glibc or libc6) before 2.20 does not properly reuse file descriptors, which allows remote attackers to send DNS queries to unintended locations via a large number of request that trigger a call to the getaddrinfo function.
CVE-2014-0475	Multiple directory traversal vulnerabilities in GNU C Library (aka glibc or libc6) before 2.20 allow context-dependent attackers to bypass ForceCommand restrictions and possibly have other unspecified impact via a .. (dot dot) in a (1) LC_*, (2) LANG, or other locale environment variable.
CVE-2014-3570	The BN_sqr implementation in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k does not properly calculate the square of a BIGNUM value, which might make it easier for remote attackers to defeat cryptographic protection mechanisms via unspecified vectors, related to crypto/bn/asm/mips.pl, crypto/bn/asm/x86_64-gcc.c, and crypto/bn/bn_asm.c.
CVE-2014-3571	OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via a crafted DTLS message that is processed with a different read operation for the handshake header than for the handshake body, related to the dtls1_get_record function in d1_pkt.c and the ssl3_read_n function in s3_pkt.c.
CVE-2014-3572	The ssl3_get_key_exchange function in s3_clnt.c in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote SSL servers to conduct ECDHE-to-ECDH downgrade attacks and trigger a loss of forward secrecy by omitting the ServerKeyExchange message.
CVE-2014-7817	The wordexp function in GNU C Library (aka glibc) 2.21 does not enforce the WRDE_NOCMD flag, which allows context-dependent attackers to execute arbitrary commands, as demonstrated by input containing "$((`...`))".
CVE-2014-8176	The dtls1_clear_queues function in ssl/d1_lib.c in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h frees data structures without considering that application data can arrive between a ChangeCipherSpec message and a Finished message, which allows remote DTLS peers to cause a denial of service (memory corruption and application crash) or possibly have unspecified other impact via unexpected application data.
CVE-2014-8275	OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k does not enforce certain constraints on certificate data, which allows remote attackers to defeat a fingerprint-based certificate-blacklist protection mechanism by including crafted data within a certificate's unsigned portion, related to crypto/asn1/a_verify.c, crypto/dsa/dsa_asn1.c, crypto/ecdsa/ecs_vrf.c, and crypto/x509/x_all.c.
CVE-2014-9297	This candidate has been reserved by an organization or individual that will use it when announcing a new security problem. When the candidate has been publicized, the details for this candidate will be provided

Table 9 - List of Security Bug Fixes

CVE	Description
CVE-2015-0204	The ssl3_get_key_exchange function in s3_clnt.c in OpenSSL before 0.9.8zd, 1.0.0 before 1.0.0p, and 1.0.1 before 1.0.1k allows remote SSL servers to conduct RSA-to-EXPORT_RSA downgrade attacks and facilitate brute-force decryption by offering a weak ephemeral RSA key in a noncompliant role, related to the "FREAK" issue. NOTE: the scope of this CVE is only client code based on OpenSSL, not EXPORT_RSA issues associated with servers or other TLS implementations.
CVE-2015-0205	The ssl3_get_cert_verify function in s3_srvr.c in OpenSSL 1.0.0 before 1.0.0p and 1.0.1 before 1.0.1k accepts client authentication with a Diffie-Hellman (DH) certificate without requiring a CertificateVerify message, which allows remote attackers to obtain access without knowledge of a private key via crafted TLS Handshake Protocol traffic to a server that recognizes a Certification Authority with DH support.
CVE-2015-0206	Memory leak in the dtls1_buffer_record function in d1_pkt.c in OpenSSL 1.0.0 before 1.0.0p and 1.0.1 before 1.0.1k allows remote attackers to cause a denial of service (memory consumption) by sending many duplicate records for the next epoch, leading to failure of replay detection.
CVE-2015-0209	Use-after-free vulnerability in the d2i_ECPrivateKey function in crypto/ec/ec_asn1.c in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a might allow remote attackers to cause a denial of service (memory corruption and application crash) or possibly have unspecified other impact via a malformed Elliptic Curve (EC) private-key file that is improperly handled during import.
CVE-2015-0285	The ssl3_client_hello function in s3_clnt.c in OpenSSL 1.0.2 before 1.0.2a does not ensure that the PRNG is seeded before proceeding with a handshake, which makes it easier for remote attackers to defeat cryptographic protection mechanisms by sniffing the network and then conducting a brute-force attack.
CVE-2015-0286	The ASN1_TYPE_cmp function in crypto/asn1/a_type.c in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not properly perform boolean-type comparisons, which allows remote attackers to cause a denial of service (invalid read operation and application crash) via a crafted X.509 certificate to an endpoint that uses the certificate-verification feature.
CVE-2015-0287	The ASN1_item_ex_d2i function in crypto/asn1/tasn_dec.c in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not reinitialize CHOICE and ADB data structures, which might allow attackers to cause a denial of service (invalid write operation and memory corruption) by leveraging an application that relies on ASN.1 structure reuse.
CVE-2015-0288	The X509_to_X509_REQ function in crypto/x509/x509_req.c in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a might allow attackers to cause a denial of service (NULL pointer dereference and application crash) via an invalid certificate key.
CVE-2015-0289	The PKCS#7 implementation in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a does not properly handle a lack of outer ContentInfo, which allows attackers to cause a denial of service (NULL pointer dereference and application crash) by leveraging an application that processes arbitrary PKCS#7 data and providing malformed data with ASN.1 encoding, related to crypto/pkcs7/pk7_doit.c and crypto/pkcs7/pk7_lib.c.

Table 9 - List of Security Bug Fixes

CVE	Description
CVE-2015-0292	Integer underflow in the EVP_DecodeUpdate function in crypto/evp/encode.c in the base64-decoding implementation in OpenSSL before 0.9.8za, 1.0.0 before 1.0.0m, and 1.0.1 before 1.0.1h allows remote attackers to cause a denial of service (memory corruption) or possibly have unspecified other impact via crafted base64 data that triggers a buffer overflow.
CVE-2015-0293	The SSLv2 implementation in OpenSSL before 0.9.8zf, 1.0.0 before 1.0.0r, 1.0.1 before 1.0.1m, and 1.0.2 before 1.0.2a allows remote attackers to cause a denial of service (s2_lib.c assertion failure and daemon exit) via a crafted CLIENT-MASTER-KEY message.
CVE-2015-1789	The X509_cmp_time function in crypto/x509/x509_vfy.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (out-of-bounds read and application crash) via a crafted length field in ASN1_TIME data, as demonstrated by an attack against a server that supports client authentication with a custom verification callback.
CVE-2015-1790	The PKCS7_dataDecodefunction in crypto/pkcs7/pk7_doit.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (NULL pointer dereference and application crash) via a PKCS#7 blob that uses ASN.1 encoding and lacks inner EncryptedContent data.
CVE-2015-1791	Race condition in the ssl3_get_new_session_ticket function in ssl/s3_clnt.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b, when used for a multi-threaded client, allows remote attackers to cause a denial of service (double free and application crash) or possibly have unspecified other impact by providing a NewSessionTicket during an attempt to reuse a ticket that had been obtained earlier.
CVE-2015-1792	The do_free_upto function in crypto/cms/cms_smime.c in OpenSSL before 0.9.8zg, 1.0.0 before 1.0.0s, 1.0.1 before 1.0.1n, and 1.0.2 before 1.0.2b allows remote attackers to cause a denial of service (infinite loop) via vectors that trigger a NULL value of a BIO data structure, as demonstrated by an unrecognized X.660 OID for a hash function.
CVE-2015-1798	The symmetric-key feature in the receive function in ntp_proto.c in ntpd in NTP 4.x before 4.2.8p2 requires a correct MAC only if the MAC field has a nonzero length, which makes it easier for man-in-the-middle attackers to spoof packets by omitting the MAC.
CVE-2015-1799	The symmetric-key feature in the receive function in ntp_proto.c in ntpd in NTP 3.x and 4.x before 4.2.8p2 performs state-variable updates upon receiving certain invalid packets, which makes it easier for man-in-the-middle attackers to cause a denial of service (synchronization loss) by spoofing the source IP address of a peer.
CVE-2015-3456	The Floppy Disk Controller (FDC) in QEMU, as used in Xen 4.5.x and earlier and KVM, allows local guest users to cause a denial of service (out-of-bounds write and guest crash) or possibly execute arbitrary code via the (1) FD_CMD_READ_ID, (2) FD_CMD_DRIVE_SPECIFICATION_COMMAND, or other unspecified commands, aka VENOM. Though the VENOM vulnerability is also agnostic of the guest operating system, an attacker (or an attacker's malware) would need to have administrative or root privileges in the guest operating system in order to exploit VENOM.

Table 9 - List of Security Bug Fixes

CVE	Description
CVE-2015-4000	The TLS protocol 1.2 and earlier, when a DHE_EXPORT ciphersuite is enabled on a server but not on a client, does not properly convey a DHE_EXPORT choice, which allows man-in-the-middle attackers to conduct cipher-downgrade attacks by rewriting a ClientHello with DHE replaced by DHE_EXPORT and then rewriting a ServerHello with DHE_EXPORT replaced by DHE, aka the “Logjam” issue.
CVE-2015-5119	This candidate has been reserved by an organization or individual that will use it when announcing a new security problem. When the candidate has been publicized, the details for this candidate will be provided.

6 Submitting a Service Request

The Mellanox® Support Center is at your service for any issues. You may access the Warranty Service through the Web Request Form by using the following link:

http://www.mellanox.com/content/pages.php?pg=support_index.